

Collaborative Learning of Anomalies with Privacy (CLAP) for Unsupervised Video Anomaly Detection: A New Baseline

Anas Al-lahham Muhammad Zaigham Zaheer Nurbek Tastan Karthik Nandakumar
Mohamed bin Zayed University of Artificial Intelligence (MBZUAI)
Abu Dhabi, UAE

{anas.al-lahham, zaigham.zaheer, nurbek.tastan, karthik.nandakumar}@mbzuai.ac.ae

Abstract

Unsupervised (US) video anomaly detection (VAD) in surveillance applications is gaining more popularity recently due to its practical real-world applications. As surveillance videos are privacy sensitive and the availability of large-scale video data may enable better US-VAD systems, collaborative learning can be highly rewarding in this setting. However, due to the extremely challenging nature of the US-VAD task, where learning is carried out without any annotations, privacy-preserving collaborative learning of US-VAD systems has not been studied yet. In this paper, we propose a new baseline for anomaly detection capable of localizing anomalous events in complex surveillance videos in a fully unsupervised fashion without any labels on a privacy-preserving participant-based distributed training configuration. Additionally, we propose three new evaluation protocols to benchmark anomaly detection approaches on various scenarios of collaborations and data availability. Based on these protocols, we modify existing VAD datasets to extensively evaluate our approach as well as existing US SOTA methods on two large-scale datasets including UCF-Crime and XD-Violence. All proposed evaluation protocols, dataset splits, and codes are available here: <https://github.com/AnasEmad11/CLAP>.

1. Introduction

Recent years have seen a surge in federated learning based methods, where the goal is to enable collaborative training of machine learning models without transferring any training data to a central server. This direction of research in machine learning is of notable importance as it enables learning with multiple participants that can contribute data without compromising privacy. Several researchers have studied federated learning for different applications such as medical diagnosis [1, 3, 6, 22], network security [10, 19, 21, 30], and large-scale classification models [4, 14, 42].

Anomaly detection in surveillance videos, being one

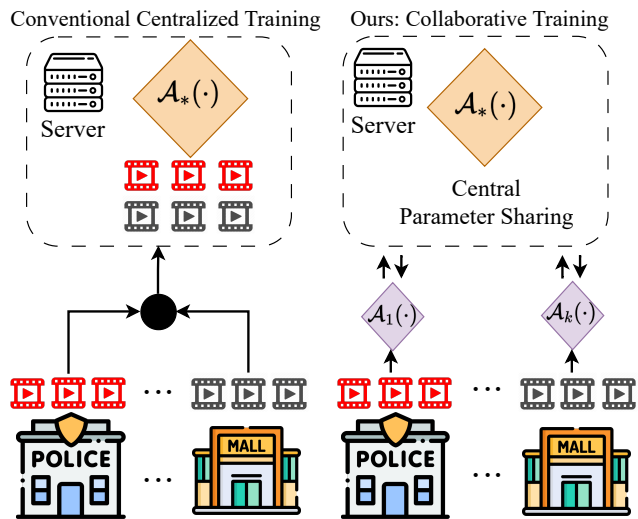


Figure 1. a) Conventional central training requires all training data to be on the server to carry out the training. This setting cannot ensure privacy, thus hindering collaborations between different entities holding large-scale surveillance data. b) Our proposed unsupervised video anomaly detection technique does not require the transfer of training data between the server and participants, thus ensuring complete privacy.

of the large-scale applications of computer vision, may greatly benefit from autonomous collaborative training algorithms. VAD in surveillance videos is privacy sensitive and may involve data belonging to several organizations/establishments. This may result in hectic bureaucratic processes to obtain data from each establishment for centralized training. For example, the police department of a city may not be willing to share street surveillance videos due to privacy concerns of the general public, or a daycare facility may have to obtain the consent of all parents to be allowed to share its CCTV footage. Such restrictions may hinder the possibility of obtaining large-scale data to train efficient anomaly detectors making a central training requiring all training data the least preferred option in the real-world scenarios. Unfortunately, to the best of our knowl-

edge, there are hardly any notable attempts to leverage federated learning for video anomaly detection which may be due to the challenging nature of the anomaly detection task itself. Anomalies are often unknown and it is not feasible to collect all possible anomaly examples for a model to learn from. Furthermore, anomalies are rare in nature, and annotating large amounts of data is laborious.

In this work, we explore video anomaly detection (VAD) on two fronts: 1) Unsupervised - videos are used without any labels. 2) Distributed participant based learning - the server does not get any raw data from participants. Unsupervised VAD is a relatively recent development in the field of anomaly detection in which no supervision is provided during training [2, 39]. This class of VAD training is different from the existing one-class classification (OCC) and weakly supervised (WS) learning. In OCC, only normal videos are provided for training whereas, in WS both normal and anomalous videos are provided with binary video-level labels. Unsupervised VAD is somewhat closer to WS-VAD, as it also utilizes large sets of videos containing normal and anomalous events. However, instead of relying on video-level labels, the network is designed in such a way that it utilizes several cues, such as the abundance of normal events/scarcity of anomalies, etc., present in the surveillance videos to drive the overall training [39]. Unsupervised VAD itself is a challenging task and the complexity increases multifold when we consider distributed participants setting for training. However, this is also more rewarding due to zero annotation labor and a more practical real-world application enabling collaboration between different large-scale data sources.

To this end, we propose CLAP, an approach for Collaborative Learning of Anomalies with Privacy that takes unlabelled videos at multiple nodes (participants) as input and collaboratively learns to predict frame-level anomaly score predictions as output (Figure 1). At an abstract level, our approach can be divided into three distinct steps: Common knowledge based data segregation for local training, knowledge accumulation at server, and local feedback. As we approach the task as fully unsupervised, i.e., without any labels, videos at each participant's end are segregated to separate normal and anomalous candidates. To this end, we propose to utilize Von Neumann entropy as a metric [25], and apply Gaussian Mixture Model (GMM) to create clusters of normal and anomalous videos. After certain local epochs, the weights of all local trainings are accumulated at the server. Based on this, a feedback loop is formed to refine the initial labels obtained by each participant and share the commonly learned knowledge between each participant.

As video anomaly detection in distributed participant settings is not well-studied, we explore multiple training and evaluation scenarios mimicking real-world collabora-

tions. These scenarios, listed in the order of complexity, include all participants having access to similar type of training data, different participants having access to different types of anomalies, and different participants having totally different types and numbers of videos. Overall, the contributions of our manuscript are: 1) We propose, CLAP, a new baseline for anomaly detection capable of localizing anomalous events in complex surveillance scenarios in a fully unsupervised fashion without any labels on a privacy-preserving participant-based training configuration. To the best of our knowledge, CLAP is the first rigorous attempt to tackle video anomaly detection in the federated learning setting. 2) We propose three new evaluation scenarios to extensively evaluate CLAP on various scenarios of collaborations and data availability. 3) To carry out these evaluations, we modify the existing VAD datasets to create new splits.

2. Related Work

2.1. Federated Learning

Federated learning has been studied for various computer vision applications including healthcare [19, 21, 30], surveillance [5, 7, 18, 24, 28], and autonomous driving [8, 9, 17, 23, 41]. Video anomaly detection has not been well-studied in federated learning settings. The closely related anomaly detection in federated learning setting mostly includes network security related methods in which different network attacks are identified from the normal traffic of packets [10, 19, 21, 30]. However, given that these are mostly supervised tasks in the form of one-class classification, the problem of distributed learning transforms into weight-sharing optimization. Recently, Doshi *et al.* [7] have proposed a weakly supervised federated learning (FL) video anomaly detection (VAD) method. The idea is to explore the FL setting by randomly dividing the data between multiple clients. In essence, this work is related to our approach as we also study FL for VAD. However, we primarily explore the unsupervised setting of VAD. Moreover, without any training labels for VAD, we additionally propose common knowledge-based data segregation for local training and local feedback for improved pseudo-labeling to carry out the collaborative training. Furthermore, we propose realistic scenarios to evaluate VAD methods in FL setting.

2.2. Unsupervised Anomaly Detection

Introduced by Zaheer *et al.* [39], fully unsupervised anomaly detection is a relatively new idea and the methods that do not require any training labels are still quite sparse in the literature. This problem is extremely challenging due to the rarity of anomalies and the complete lack of supervision labels. Zaheer *et al.* [39], by relying on the abundance of normal data, proposed to first train a generative model to learn the overall normal trends in the dataset. A

classifier model is then trained based on the pseudo labels obtained from the generator. Both models are then trained in a cooperative manner to converge as an anomaly detector. This line of research has been extended by Anas *et al.* [2]. They have proposed to utilize hierarchical clustering to obtain fine-grained pseudo-labels. The training of an anomaly detector is then carried out using these pseudo labels. Our approach also begins with pseudo-label generation, followed by training and then a feedback loop to improve the pseudo-labels. However, we attempt to address the problem in collaborative learning setting where privacy-preserving training is carried out by multiple participants. In addition, we propose a common knowledge aware data segregation in which all participants share common clusters knowledge to obtain pseudo-labels for training.

We also acknowledge one-class classification (OCC) [11, 12, 27, 35] and weakly-supervised (WS) [29, 31, 34, 36, 37, 40] approaches for video anomaly detection. To enrich the quality of analysis and to provide several aspects of our research work, we carry out some additional experiments on weakly-supervised settings by incorporating weak labels in our training in Section 4. However, the prime focus of our research work is unsupervised video anomaly detection, and thus, OCC and WS anomaly detection are not in the scope of our research work.

3. Methodology

Problem Definition: Given a dataset of training videos without any labels, the goal of US-VAD is to learn an anomaly detector $\mathcal{A}_\theta(\cdot)$ that classifies each frame in a given test video V_* as either *normal* (0) or *anomalous* (1). Suppose that there are K participants $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_K$ for collaborative training and each participant \mathcal{P}_k has its own local training dataset $\mathcal{D}_k = \{V_{1,k}, V_{2,k}, \dots, V_{N_k,k}\}$ containing N_k videos used to train its own local anomaly detector model $\mathcal{A}_{\theta_k}(\cdot)$, $k \in [1, K]$. It is assumed that all participants share a common test set and the performance of the model $\mathcal{A}_{\theta_k}(\cdot)$ on this test set is denoted by \mathcal{O}_k . The goal of each participant is to collaborate with other participants in order to obtain a global model θ_* that has better performance \mathcal{O}_* on the test set compared to all \mathcal{O}_k , without compromising the privacy of participant’s local data \mathcal{D}_k .

Preprocessing: For simplicity of notation, we drop the participant index k unless required. Let the training dataset of N videos at a generic participant be denoted as $\mathcal{D} = \{V_1, V_2, \dots, V_N\}$. We split each video V_i into a sequence of m_i non-overlapping segments S_{ij} , where each segment is in turn composed of r frames. Note that $i \in [1, N]$ refers to the video index and $j \in [1, m_i]$ is the segment index within a video. Unlike many state-of-the-art AD methods [26, 29, 32, 33] that compress each video into a fixed number of segments (i.e., $m_i = m, \forall i \in [1, N]$) along the temporal axis, we avoid any compression and make use of all

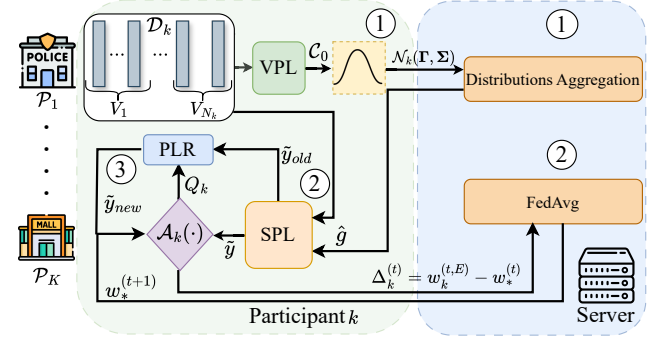


Figure 2. Architecture of CLAP, an unsupervised video anomaly detection model trained by multiple collaborating participants.

available non-overlapping segments¹. For each segment S_{ij} , a feature vector $\mathbf{f}_{ij} \in \mathbb{R}^d$ is obtained using a pre-trained feature extractor $\mathcal{F}(\cdot)$.

High-level Overview: Our proposed CLAP model for collaborative training consists of three main stages. The aim is to generate both video-level and segment-level pseudo-labels to enable the training of the anomaly detector model $\mathcal{A}_\theta(\cdot)$. In the first Common Knowledge-based Data Segregation (CKDS) stage, the generation of segment-level pseudo labels is done in a collaborative manner. We generate a video-level pseudo-label $\hat{y}_i \in \{0, 1\}$, $i \in [1, N]$ for each video in the training set using hierarchical divisive clustering. Subsequently, segment-level pseudo-labels $\tilde{y}_{ij} \in \{0, 1\}$, $i \in [1, N]$, $j \in [1, m_i]$ are generated for all the segments in the training set through collaborative statistical hypothesis testing. Finally, we train a local anomaly detector $\mathcal{A}_\theta(\cdot) : \mathbb{R}^d \rightarrow [0, 1]$ that assigns an anomaly score between 0 and 1 (higher values indicate higher confidence of being an anomaly) to the given video segment based on its feature representation \mathbf{f}_{ij} .

During training, we utilize both server knowledge accumulation (SKA) and local feedback stages to improve the performance of the local model. In the second (SKA) stage, we use the well-known Federated Averaging (FedAvg) algorithm [20] (we also analyze other FL aggregation methods in Section 6 of supplementary material) for aggregating local anomaly detection models. Finally, upon completion of a pre-determined number of collaboration rounds to update the weights given the initial pseudo-labels, our algorithm initiates the local feedback or pseudo-label refinement (PLR) stage. During this stage, we use confidence scores predicted by the network to refine the generated pseudo-labels from the first stage.

¹Note that for simplicity, we use the notation val_{ij} to represent a value for segment j in video V_i , val_i represents the set of values of all segments in video V_i , and val simply represents the collection of all segment-wise values of all the videos in the dataset \mathcal{D} .

3.1. Knowledge-based Data Segregation (CKDS)

At the participant level, since the training does not assume any labels, we first generate pseudo-labels for the videos in the training set by clustering them into two groups: normal and anomalous (see Alg. 1).

Video-Level Pseudo-Labels: Previous works in WS-VAD have shown that normal video segments have lower temporal feature magnitude compared to anomalous segments [31]. In addition, we observe the variance of the difference in the feature magnitude between consecutive segments in a given anomalous video is higher than in a normal video. Furthermore, we consider von-Neumann entropy H of the covariance matrix computed based on the features as an indicator of the presence of anomalies, i.e., the entropy of segments is generally expected to be lower for normal videos. Based on these cues, we represent each video V_i using a statistical summary $\mathbf{x}_i = [\sigma_i, H_i]$ of its features as follows:

$$\mu_i = \frac{1}{(m_i - 1)} \sum_{j=1}^{(m_i-1)} (\|\mathbf{f}_{ij}\|_2 - \|\mathbf{f}_{i(j+1)}\|_2), \quad (1)$$

$$\sigma_i = \sqrt{\frac{1}{(m_i - 2)} \sum_{j=1}^{(m_i-1)} ((\|\mathbf{f}_{ij}\|_2 - \|\mathbf{f}_{i(j+1)}\|_2) - \mu_i)^2}, \quad (2)$$

$$\text{Cov}[\mathbf{f}_{i,1}, \dots, \mathbf{f}_{i,m_i}] = \mathbf{U}_i \boldsymbol{\Sigma}_i \mathbf{V}_i^T \quad (3)$$

$$H_i = -\text{tr}[\boldsymbol{\Sigma}_i \log \boldsymbol{\Sigma}_i] \quad (4)$$

where $\|\cdot\|_2$ represents the ℓ_2 norm of a vector. Thus, each video V_i is represented using a 2D vector \mathbf{x}_i , corresponding to the variance σ_i and entropy H_i of the video segments. Videos in the training set are then divided into two clusters (\mathcal{C}_{s_1} and \mathcal{C}_{s_2}), with $|\mathcal{C}_{s_1}|$ and $|\mathcal{C}_{s_2}|$ samples, respectively, based on the above representation \mathbf{x}_i .

Intuitively, building on the assumption that normal samples have a lower entropy, we compute the average entropy of samples in each cluster and assign the cluster with the larger average entropy as anomalous (0), while the other cluster is labeled as normal (1). At the end of this stage, all the videos in the training set are assigned a pseudo-label based on their corresponding cluster label, i.e., $\hat{y}_i = s$, if $\mathbf{x}_i \in \mathcal{C}_s$, where $s \in \{0, 1\}$.

Segment-Level Pseudo-Labels: All the segments from videos that are ‘‘pseudo-labeled’’ as normal ($\hat{y}_i = 0$) by the previous stage can be considered normal. However, most of the segments in an anomalous video are also normal due to the smaller temporal extent of anomalies. To tackle this, we treat the detection of anomalous segments as a statistical hypothesis-testing problem. Specifically, the null hypothesis is that a given video segment is normal. By modeling

Algorithm 1 Video-level Pseudo-Label Generation (VPL)

Input: dataset $\mathcal{D} = \{V_1, \dots, V_N\}$, feature extractor $\mathcal{F}(\cdot)$

- 1: **for** $i = 1$ to N **do**
- 2: Partition V_i into m_i segments $[S_{i1}, \dots, S_{im_i}]$
- 3: Extract segment features $[\mathbf{f}_{i1}, \dots, \mathbf{f}_{im_i}]$ using $\mathcal{F}(\cdot)$
- 4: Compute $\mathbf{x}_i = [\sigma_i, H_i]$ using Eqs. 2 & 4
- 5: **end for**
- 6: $\mathcal{C}_{s_1}, \mathcal{C}_{s_2} \leftarrow \text{Clustering}(\text{GMM})$, where $a = |\mathcal{C}_{s_1}|, b = |\mathcal{C}_{s_2}|$
- 7: **if** $\frac{1}{a} \sum_{i=1}^a H_i > \frac{1}{b} \sum_{i=1}^b H_i$ **then**
- 8: $\mathcal{C}_0 = \mathcal{C}_{s_2}, \mathcal{C}_1 = \mathcal{C}_{s_1}$
- 9: **else**
- 10: $\mathcal{C}_0 = \mathcal{C}_{s_1}, \mathcal{C}_1 = \mathcal{C}_{s_2}$
- 11: **end if**
- 12: $\forall i \in [1, N], \hat{y}_i \leftarrow 0$ **if** $\mathbf{x}_i \in \mathcal{C}_0$, **else** $\hat{y}_i \leftarrow 1$

return $\hat{\mathbf{y}} = \{\hat{y}_1, \dots, \hat{y}_N\}$

the distribution of features under the null hypothesis as a Gaussian distribution, we identify the anomalous segments by estimating their p-value and rejecting the null hypothesis if the p-value is less than the significance level α .

To model the distribution of features (**at the participant level**) under the null hypothesis, we consider only the segments from videos that are pseudo-labeled as normal by the VPL stage (see Algo. 1). Let $\mathbf{z}_{ij} \in \mathbb{R}^d$ be a low-dimensional representation of a segment S_{ij} . In this work, we simply set $\mathbf{z}_{ij} = \|\mathbf{f}_{ij}\|_2$. We assume that \mathbf{z}_{ij} follows a Gaussian distribution $\mathcal{N}(\boldsymbol{\Gamma}, \boldsymbol{\Theta})$ under the null hypothesis and estimate the parameters $\boldsymbol{\Gamma}$ and $\boldsymbol{\Theta}$ as follows:

$$\boldsymbol{\Gamma} = \frac{1}{M_0} \sum_{i=1, \hat{y}_i=0}^N \sum_{j=1}^{m_i} \mathbf{z}_{ij}, \quad (5)$$

$$\boldsymbol{\Theta} = \frac{1}{(M_0 - 1)} \sum_{i=1, \hat{y}_i=0}^N \sum_{j=1}^{m_i} (\mathbf{z}_{ij} - \boldsymbol{\Gamma})(\mathbf{z}_{ij} - \boldsymbol{\Gamma})^T, \quad (6)$$

where $M_0 = \sum_{i=1, \hat{y}_i=0}^N m_i$.

Let $(\boldsymbol{\Gamma}_k, \boldsymbol{\Theta}_k)$ be the Gaussian parameters at participant \mathcal{P}_k based on $M_{0,k}$ normal segments. The participants share these parameters with the server and the server sends back a Gaussian mixture model \mathcal{G} . One simple way to construct \mathcal{G} is to treat $(\boldsymbol{\Gamma}_k, \boldsymbol{\Theta}_k)$ of each participant as a mixture component and weight these components based on the corresponding number of normal segments. More sophisticated aggregation approaches could also be employed by the server. Subsequently, for all the segments in videos that are pseudo-labeled as anomalous, the p-value is computed as:

$$p_{ij} = \mathcal{G}(\mathbf{z}_{ij}), \quad \forall \hat{y}_i = 1. \quad (7)$$

If $p_{ij} < \alpha$, the segment can be directly assigned a pseudo-label of 1. However, we identify a potential anomalous region by sliding a window of size w_i across the video and se-

lecting the region that has the lowest average p-values (i.e., $\min_l \left\{ \frac{1}{w_i} \sum_{j=l+1}^{l+w_i} p_{ij}, \forall l \in [0, m_i - w_i] \right\}$). Each segment present in this anomalous region is assigned a pseudo-label of 1, while all the remaining segments are pseudo-labeled as normal with a value of 0. Thus, a pseudo-label $\tilde{y}_{ij} \in \{0, 1\}$ is assigned to all the segments in the training set. This window-based labeling may be seen as utilizing the temporal consistency property of the surveillance videos commonly utilized in the existing literature [29, 36].

Algorithm 2 Segment-level Pseudo-Label Gen. (SPL)

Input: Mode, $0 < \beta \leq 1$, \hat{y} (video-level pseudo-labels) and Gaussian mixture model \mathcal{G} for Generate mode, \tilde{y} (current segment-level pseudo-labels) and Q (segment-level confidence scores) for Update mode

Mode I: Generate

```

1: for  $i = 1$  to  $N$  do
2:   if  $\hat{y}_i = 1$  then
3:     Compute  $p_{ij}$  using Eq. 7,  $\forall j \in [1, m_i]$ 
4:      $w_i \leftarrow \lceil \beta m_i \rceil$ 
5:      $l_i = \arg \min_l \left\{ \frac{1}{w_i} \sum_{j=l+1}^{l+w_i} p_{ij}, \forall l \in [0, m_i - w_i] \right\}$ 
6:      $\tilde{y}_{ij} \leftarrow 1, \forall j \in [l_i + 1, l_i + w]$ 
7:   end if
8: end for
return  $\tilde{y}$ 

```

Mode II: Update (PLR)

```

9: for  $i = 1$  to  $N$  do
10:  Set  $q_{ij}$  based on  $Q \forall j \in [1, m_i]$ 
11:   $w_i \leftarrow \lceil \beta m_i \rceil$ 
12:   $l_i = \arg \max_l \left\{ \frac{1}{w_i} \sum_{j=l+1}^{l+w_i} q_{ij}, \forall l \in [0, m_i - w_i] \right\}$ 
13:   $\tilde{q}_{ij} \leftarrow 0, \forall j \in [1, m_i]$ 
14:   $\tilde{q}_{ij} \leftarrow 1, \forall j \in [l_i + 1, l_i + w]$ 
15:   $\hat{q}_{ij} \leftarrow 0, \forall j \in [1, m_i]$ 
16:  if  $\text{len}(\tilde{y}_i \cap \tilde{q}_i) > 0$  then
17:     $\hat{q}_{ij} \leftarrow 1, \forall j \in [1, m_i], \text{if } \tilde{q}_{ij} \in (\tilde{y}_i \cap \tilde{q}_i)$ 
18:  else
19:     $\hat{q}_{ij} \leftarrow 1, \forall j \in [1, m_i], \text{if } \tilde{q}_{ij} \neq 0 \text{ or } \tilde{y}_{i,j} \neq 0$ 
20:  end if
21: end for
22: return  $\hat{q}$ 

```

3.2. Server Knowledge Accumulation (SKA) and Local Feedback

At the beginning of this stage, each participant would have obtained segment-level pseudo-labels for its own dataset by sequentially applying Algo. 1 and Algo. 2 as described in Section 3.1. As earlier, if \mathcal{D} is the unlabeled training dataset of a generic participant, we can obtain the segment-level pseudo-labeled training set $\tilde{\mathcal{D}} = \{(\mathbf{f}_{ij}, \tilde{y}_{ij})\}$ containing M samples, where $i \in [1, N]$, $j \in [1, m_i]$, and $M = \sum_{i=1}^n m_i$. This labeled training set $\tilde{\mathcal{D}}$ can be used to train the anomaly

Algorithm 3 CLAP

Require: Local training dataset \mathcal{D}_k . Server initializes parameter $\theta_*^{(0)}$.

```

1: for each participant  $k = 1, 2, \dots, K$  do
2:    $\hat{y}_k \leftarrow$  Algorithm 1 ( $\mathcal{D}_k$ )
3:    $\forall i \in [1, N_k], j \in [1, m_i], \tilde{y}_{ij} \leftarrow 0$ , Compute  $\mathbf{z}_{ij}$ 
4:   Compute  $(\Gamma_k, \Theta_k)$  using Eqs. 5 & 6
5:   return  $(\Gamma_k, \Theta_k)$  to server
6: end for
7: server:  $\mathcal{G} \leftarrow$  Mixture of Gaussians ( $\{(\Gamma_k, \Theta_k)\}_{k=1}^K$ )
8: for each participant  $k = 1, 2, \dots, K$  do
9:    $\tilde{y}_k \leftarrow$  Algorithm 2 (Generate,  $\hat{y}_k, \mathcal{G}$ )
10: end for
11: for each round  $t = 0, 1, \dots, T$  do
12:   for each participant  $k = 1, 2, \dots, K$  do
13:      $\tilde{\mathcal{D}}_k \leftarrow (\mathcal{D}_k, \tilde{y}_k)$ 
14:      $\theta_k^{(t,0)} \leftarrow \theta_*^{(t)}$ 
15:     for local iteration  $e = 0, 1, \dots, E$  do
16:        $\theta_k^{(t,e+1)} \leftarrow \theta_k^{(t,e)} - \eta \cdot \nabla \mathcal{L}_{total,k}(\tilde{\mathcal{D}}_k, \theta_k^{(t,e)})$ 
17:     end for
18:      $\Delta_k^{(t)} \leftarrow \theta_k^{(t,E)} - \theta_*^{(t)}$ 
19:      $Q_k \leftarrow \mathcal{A}_{\theta_k^{(t,E)}}(\mathcal{D}_k)$ 
20:      $\tilde{y}_k \leftarrow$  Algorithm 2 (Update,  $Q_k, \tilde{y}_k$ )
21:     return  $\Delta_k^{(t)}$  to server
22:   end for
23:   server:  $\theta_*^{(t+1)} \leftarrow \theta_*^{(t)} + \frac{\lambda}{K} \sum_{k=1}^K \Delta_k^{(t)}$ 
24: end for

```

detector $\mathcal{A}_\theta(\cdot)$ by minimizing the following objective:

$$\min_{\theta} \mathcal{L}_{total} = \sum_{i=1}^N \sum_{j=1}^{m_i} \mathcal{L}(\mathcal{A}_\theta(\mathbf{f}_{ij}), \tilde{y}_{ij}), \quad (8)$$

where \mathcal{L} is an appropriate loss function and θ denotes the parameters of the anomaly detector $\hat{\mathcal{A}}(\cdot)$. Stochastic gradient descent (SGD) is used for the above optimization. Following recent state-of-the-art methods [32, 36, 39], a basic neural network architecture is considered for our anomaly detector (see Supplementary for more details).

Now, we proceed to describe the collaborative training of the anomaly detector, which is referred to as **server knowledge accumulation** (SKA). At the beginning of each collaboration round t , the server broadcasts the current global model parameters $\theta_*^{(t)}$ to all the participants. Using these global parameters as the initialization, each participant will perform E local SGD iterations to get the updated parameters $\theta_k^{(t,E)}$. At the end of the local training, the participant sends the local gradient $\Delta_k^{(t)} \leftarrow \theta_k^{(t,E)} - \theta_*^{(t)}$ back to the server. The server aggregates these gradients and applies the update to the global model as $\theta_*^{(t+1)} \leftarrow \theta_*^{(t)} + \frac{\lambda}{K} \sum_{i=1}^K \Delta_k^{(t)}$ as shown in Algo. 3, where λ is the learning rate.

In addition to SKA, we also incorporate local feedback or **pseudo-label refinement process** (PLR) as shown in Algo.2(Update). During this stage, we use confidence scores Q predicted by the local model to refine the segment-level pseudo-labels. The aim is to use the high-confidence segments to update the pseudo-labels \tilde{y} generated from Algo.2(Generate). First, we determine the maximum confidence region by a sliding window w_i similar to Algo.2(Generate) and assign those segments as $\tilde{q}_{ij} = 1$. The refinement of the old pseudo-labels is based on two rules. First, if there is an intersection between the maximum confidence region and the generated pseudo-labels ($\tilde{y}_i \cap \tilde{q}_i$) > 0 , we assign all the segments in \tilde{q}_{ij} that are in the intersection set a value of 1. On the other hand, if there is no intersection, we assume that the old pseudo-labels missed an additional anomalous window. Therefore, the whole of the maximum confidence region will be assigned as anomalous.

3.3. Inference

At the end of all communication rounds, the final global model $\mathcal{A}_{\theta^*(x)}(\cdot)$ is used for inference. A given test video V_* is partitioned into m_* non-overlapping segments S_{*j} , $j \in [1, m_*]$. Feature vectors \mathbf{f}_{*j} are extracted from each segment using $\mathcal{F}(\cdot)$, which are directly passed to the trained detector $\mathcal{A}_{\theta^*(x)}(\cdot)$ to obtain segment-level anomaly score predictions. As the final goal is frame-level anomaly scores, all frames within a segment of the test video inherit the predicted anomaly score for that corresponding segment.

4. Experiments and Results

4.1. Datasets and Implementation

We evaluate CLAP and conduct SOTA comparisons on two publicly available large-scale datasets, UCF-Crime and XD-Violence. Due to limited space, datasets and implementation details are provided in Supplementary.

4.2. Training Settings

As the aim of CLAP is to train an anomaly detector with collaboration between multiple participants, we provide comparisons of our approach with existing SOTA anomaly detection methods under the following three different training settings:

Centralized Training: This is the conventional setting of training where privacy is not ensured and the participants have to send all training data to the server for joint training. Anomaly detection performance is measured on the complete test set.

Local Training: This setting assumes that participants are not collaborating and each individual participant trains its anomaly detector locally with its own data. Anomaly detection performance is measured for each participant individually on the complete test set.

	Method	UCF-Crime	XD-Violence
Centralized	GCL [39]	71.04	73.62
	C2FPL [2]	80.65	80.09
	CLAP	80.9	81.71
Local	GCL [39]	56.63	58.11
	C2FPL [2]	61.33	60.05
	CLAP	63.93	62.37
Collaborative	GCL [39]	67.12	68.19
	C2FPL [2]	75.20	74.36
	CLAP	78.02	77.65

Table 1. AUC performance comparisons of unsupervised SOTA on UCF-Crime and XD-Violence datasets for five participants.

Collaborative Training: In this setting, all participants collaborate to train a joint anomaly detector. Participants do not need to send their training data to the server to carry out the training. Anomaly detection performance of the jointly trained model is measured on the complete test set.

4.3. Comparisons with Unsupervised SOTA

Table 1 summarizes the results of the existing unsupervised anomaly detection approaches on the three training settings: centralized, local, and collaborative. We re-implemented the existing methods [2, 39] on local and collaborative training settings for fair and detailed comparisons. In essence, centralized training is the upper bound of the collaborative training whereas local training is the lower bound. While better performance compared to the local training setting may indicate the success of collaborative learning, a better model should also demonstrate minimal performance difference from the centralized training. CLAP demonstrates AUC performances of 80.91% and 81.71% on UCF-crime and XD-Violence datasets on the centralized setting (upper bound). In the local setting (lower bound) with five participants, CLAP demonstrate AUC performances of 63.93% and 62.37% on the two datasets. In the collaborative learning setting, CLAP demonstrates AUC performances of 78.02% and 77.65%. Overall, the results of CLAP are not only better than the existing SOTA unsupervised VAD methods but are also on par with its counterpart centralized training setting.

4.4. Comparisons with Weakly-supervised SOTA

An unsupervised method can be converted into a supervised method upon the availability of labels [39]. We explore this supervision mode under centralized, local, and collaborative training settings and report the results in Table 2. Compared to the methods developed specifically for unsupervised video anomaly detection [2, 39], CLAP demonstrates consistent performance improvements when video-level labels are present. CLAP also outperforms PRV [7], a weakly-supervised approach designed specifically for collaborative learning. Compared to other centralized approaches that do not facilitate unsupervised training, CLAP demonstrates better performance than the compared meth-

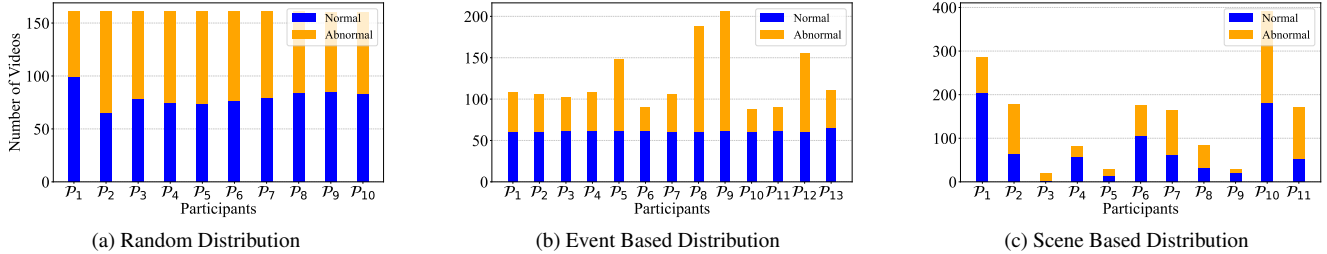


Figure 3. Distribution of UCF-Crime dataset videos based on the three training data organizations proposed in our paper to evaluate collaborative learning approaches for video Anomaly Detection. (a) Random distribution of the videos is the baseline in which each participant has an almost identical number of videos and the classes are balanced. (b) Each participant holds videos containing certain types of anomalous events such as shooting, robbery, etc. It is a relatively complex setting with the number of videos and class balance varying slightly between participants. (c) Each participant holds videos belonging to certain scenes such as shops, offices, etc. This is the most challenging setting where severe data and class imbalance are present across participants.

	Method	Unsup. Possible?	UCF-Crime	XD-Violence
Centralized	Sultani <i>et al.</i> [29]	✗	75.41	-
	RTFM [31]	✗	84.30	89.34
	MSL [15]	✗	85.30	-
	S3R [33]	✗	85.99	53.52
	CLAWS+ [38]	✗	80.90	-
	PRV [7]	✗	86.30	-
	GCL [39]	✓	79.84	82.18
	C2FPL [2]	✓	83.40	89.34
	CLAP	✓	85.50	90.04
Local	GCL [39]	✓	65.32	59.91
	C2FPL [2]	✓	65.85	63.4
	CLAP	✓	67.47	64.97
Collab.	PRV [7]	✗	82.90	-
	GCL [39]	✓	76.82	75.21
	C2FPL [2]	✓	77.60	76.98
	CLAP	✓	83.23	85.67

Table 2. AUC performance comparisons of weakly supervision SOTA on UCF-Crime and XD-Violence datasets.

ods on XD-Violence dataset and comparable performance on UCF-Crime dataset. Nevertheless, the goal of this work is not to surpass performance numbers on certain tasks but to demonstrate the possibility of unsupervised training under a collaborative learning setting to facilitate a novel research direction in the field of video anomaly detection.

4.5. Collaborative Learning in VAD: A Case Study of Different Possible Scenarios

In this section, we further explore the collaborative learning of video anomaly detection by proposing various scenarios of collaboration and consequent re-organization of the training data. Furthermore, we analyze and discuss the performance of CLAP under these scenarios.

4.5.1 Training Data Splitting

In real-world VAD applications, common sources of surveillance videos could be different government entities (e.g., department of transport or police) or private CCTV

operating institutions (e.g., shopping malls or elderly-care facilities). CLAP is designed to enable collaborative learning of a joint anomaly detector between such data sources while eliminating the need to share training data. Considering this, we propose three different training data splits mimicking different kinds of collaborations between the participants. Each of these is explained below.

Random Split: A baseline setting where each participant has randomly distributed and equal number of anomalous & normal videos for training. Figure 3a visualizes the random distribution of videos between ten participants on the UCF-crime dataset.

Event Class Based Split: In this setting, each participant has training videos based on the anomalous events present within. For example, one participant may have road accident videos whereas another participant may have robbery videos. This setting is more challenging than random distribution because each participant may have a different number of videos. Figure 3b visualizes the distribution of videos between thirteen participants on the UCF-crime dataset.

Scene Based Split: In this setting, each participant gets the videos based on the scenes/locations where the videos are recorded. For example, one participant may have surveillance videos for fuel stations, another participant may have indoor videos of offices, and so on. This is the most challenging and representative setting as the dataset is not balanced either among participants or within a participant for the normal and anomalous classes. Figure 3c visualizes the distribution of videos between eleven participants on the UCF-crime dataset.

The intuition behind these splits is that, in real-world scenarios, several participants for training a joint model may belong to different entities with different types of video data available at their disposal. For example, a police department may have videos related to street crimes, a children’s daycare facility may have surveillance available for abuse or bullying, and a mall may have videos about steal-

Split	Participants	AUC(%)
Random(IID)	5	78.02
	10	77.4
	25	76.54
	50	67.92
Event	13	77.35
Scene	11	73.99

Table 3. AUC% performance of CLAP on UCF-Crime dataset using various proposed training splits under collaborative learning setting.




Figure 4. Left: Comparison between GCL [39] and CLAP with varying number of participants. Right: Number of participants with weak video-level supervision available.

ing or shoplifting. More details on the data splitting strategies are provided in the Supplementary.

4.5.2 Experiments Using Data Splits

We conduct experiments on CLAP by splitting the UCF-crime dataset using the training data splits proposed (Figure 3) and report the results in Table 3. For the random splitting, we vary the number of participants between 5 and 50 to additionally analyze the impact of the number of participants on training. Unsurprisingly, random split based training yields the highest AUC performance of 78.02% with the participant number set to five. Experiments using event based splitting results in a closer performance of 77.35%. This demonstrates that CLAP can efficiently handle data variations and partially imbalanced data among participants. With the most challenging training setting, scene based splitting, CLAP achieves an AUC of 73.99%.

4.6. Analysis and Discussions

On Partial Weak-Supervision: The typical protocol for evaluating unsupervised methods under weakly-supervised settings is fairly simple. Once pseudo labels are generated for the whole dataset, for each video labeled as normal in the weakly supervised ground truth, label correction on pseudo labels is applied before carrying out the training [39]. In the collaborative learning setting for real-world scenarios, ensuring the availability of any form of training labels means requiring each participant to annotate their data before participating in the collaborative training. While CLAP is fully unsupervised, meaning no labels are

FedAVG	SKA	PLR	AUC(%)
✓	✗	✗	76.2
✓	✓	✗	77.1
✓	✓	✓	78.02

Table 4. Ablation study of CLAP on UCF-Crime dataset. SKA: Server Knowledge Accumulation, PLR: Pseudo Label Refinement,

required for training, in this section we explore an interesting scenario where some of the participants may have labels available for training. Figure 4 (right) shows the results of CLAP on this setting using 25 participants. CLAP demonstrates consistent performance gains when more participants contribute with video-level labels towards the collaborative training.

On Varying Number of Participants: To analyze the impact of varying numbers of participants on the unsupervised VAD training, we conduct a series of experiments using CLAP and GCL [39] with different numbers of participants and report the results in Figure 4 (left). Overall, the performance stays comparable when the participant number is set to 5, 10, and 25. However, it drops notably when the participant number is set to 50. This may be attributed to the drop is the number of videos per participant, given that the dataset size remains the same. With a more large-scale dataset, our approach may be able to accommodate an even larger number of participants without dropping performance.

Ablation: To evaluate the contribution of the various components in CLAP, we conduct an ablation study and report the results in Table 4. As seen, with each added component including SKA: Server Knowledge Accumulation stage and PLR: Pseudo-label refinement over the baseline training, notable performance gains are observed. This demonstrates the importance of all components proposed in CLAP towards unsupervised VAD.

5. Conclusion

We proposed a new baseline for anomaly detection capable of localizing anomalous events in a fully unsupervised fashion on a privacy-preserving collaborative learning configuration. We also introduced three new evaluation scenarios to extensively study anomaly detection approaches on various scenarios of collaborations and data availability. Using these scenarios, we evaluate our approach on two large-scale datasets including UCF-crime and XD-violence. A **limitation** of our approach is that the performance drops when the number of participants increases. Although some performance drop is expected in such a situation, we believe it is partly because of the limited training data available to each participant in this case. This can be addressed by curating large-scale anomaly detection datasets designed specifically for collaborative training.

References

- [1] Mohammed Adnan, Shivam Kalra, Jesse C Cresswell, Graham W Taylor, and Hamid R Tizhoosh. Federated learning and differential privacy for medical image analysis. *Scientific reports*, 12(1):1953, 2022. **1**
- [2] Anas Al-lahham, Nurbek Tastan, Zaigham Zaheer, and Karthik Nandakumar. A coarse-to-fine pseudo-labeling (c2fpl) framework for unsupervised video anomaly detection. *arXiv preprint arXiv:2310.17650*, 2023. **2, 3, 6, 7**
- [3] Faris Almalik, Naif Alkhunaizi, Ibrahim Almakky, and Karthik Nandakumar. Fesvibs: Federated split learning of vision transformer with block sampling. *arXiv preprint arXiv:2306.14638*, 2023. **1**
- [4] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kidon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, et al. Towards federated learning at scale: System design. *Proceedings of machine learning and systems*, 1:374–388, 2019. **1**
- [5] Bouziane Brik, Adlen Ksentini, and Maha Bouaziz. Federated learning for uavs-enabled wireless networks: Use cases, challenges, and open problems. *IEEE Access*, 8:53841–53849, 2020. **2**
- [6] Erfan Darzidehkalani, Mohammad Ghasemi-Rad, and PMA van Ooijen. Federated learning in medical imaging: part i: toward multicentral health care ecosystems. *Journal of the American College of Radiology*, 19(8):969–974, 2022. **1**
- [7] Keval Doshi and Yasin Yilmaz. Privacy-preserving video understanding via transformer-based federated learning. *Data Science Conference*, 2023. **2, 6, 7**
- [8] Zhaoyang Du, Celimuge Wu, Tsutomu Yoshinaga, Kok-Lim Alvin Yau, Yusheng Ji, and Jie Li. Federated learning for vehicular internet of things: Recent advances and open issues. *IEEE Open Journal of the Computer Society*, 1:45–61, 2020. **2**
- [9] Lidia Fantauzzo, Eros Fani, Debora Caldarola, Antonio Tavera, Fabio Cermelli, Marco Ciccone, and Barbara Caputo. Feddrive: Generalizing federated learning to semantic segmentation in autonomous driving. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11504–11511. IEEE, 2022. **2**
- [10] Bimal Ghimire and Danda B Rawat. Recent advances on federated learning for cybersecurity and cybersecurity for federated learning for internet of things. *IEEE Internet of Things Journal*, 9(11):8229–8249, 2022. **1, 2**
- [11] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1705–1714, 2019. **3**
- [12] Radu Tudor Ionescu, Fahad Shahbaz Khan, Mariana-Iuliana Georgescu, and Ling Shao. Object-centric auto-encoders and dummy anomalies for abnormal event detection in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7842–7851, 2019. **3**
- [13] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, pages 5132–5143. PMLR, 2020. **1**
- [14] Fan Lai, Yinwei Dai, Sanjay Singapuram, Jiachen Liu, Xiangfeng Zhu, Harsha Madhyastha, and Mosharaf Chowdhury. FedScale: Benchmarking model and system performance of federated learning at scale. In *International Conference on Machine Learning*, pages 11814–11827. PMLR, 2022. **1**
- [15] Shuo Li, Fang Liu, and Licheng Jiao. Self-training multi-sequence learning with transformer for weakly supervised video anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1395–1403, 2022. **7**
- [16] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020. **1**
- [17] Yijing Li, Xiaofeng Tao, Xuefei Zhang, Junjie Liu, and Jin Xu. Privacy-preserved federated learning for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):8423–8434, 2021. **2**
- [18] Yang Liu, Anbu Huang, Yun Luo, He Huang, Youzhi Liu, Yuanyuan Chen, Lican Feng, Tianjian Chen, Han Yu, and Qiang Yang. Fedvision: An online visual object detection platform powered by federated learning. In *Proceedings of the AAAI conference on artificial intelligence*, pages 13172–13179, 2020. **2**
- [19] Yi Liu, Jialiang Peng, Jiawen Kang, Abdullah M Iliyasu, Dusit Niyato, and Ahmed A Abd El-Latif. A secure federated learning framework for 5g networks. *IEEE Wireless Communications*, 27(4):24–31, 2020. **1, 2**
- [20] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. **3**
- [21] Viraaji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021. **1, 2**
- [22] Dianwen Ng, Xiang Lan, Melissa Min-Szu Yao, Wing P Chan, and Mengling Feng. Federated learning: a collaborative effort to achieve better medical imaging models for individual sites that have small labelled datasets. *Quantitative Imaging in Medicine and Surgery*, 11(2):852, 2021. **1**
- [23] Anh Nguyen, Tuong Do, Minh Tran, Binh X Nguyen, Chien Duong, Tu Phan, Erman Tjiputra, and Quang D Tran. Deep federated learning for autonomous driving. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 1824–1830. IEEE, 2022. **2**
- [24] Yiran Pang, Zhen Ni, and Xiangnan Zhong. Federated learning for crowd counting in smart surveillance systems. *IEEE Internet of Things Journal*, 2023. **2**
- [25] Dénes Petz. Entropy, von neumann and the von neumann entropy: Dedicated to the memory of alfred wehrl. In *John*

- von Neumann and the foundations of quantum physics, pages 83–96. Springer, 2001. [2](#)
- [26] Didik Purwanto, Yie-Tarng Chen, and Wen-Hsien Fang. Dance with self-attention: A new look of conditional random fields on anomaly detection in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 173–183, 2021. [3](#)
- [27] Mohammad Sabokrou, Mohammad Khalooei, Mahmood Fathy, and Ehsan Adeli. Adversarially learned one-class classifier for novelty detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3379–3388, 2018. [3](#)
- [28] Abdelkarim Ben Sada, Mohammed Amine Bouras, Jianhua Ma, Huang Runhe, and Huansheng Ning. A distributed video analytics architecture based on edge-computing and federated learning. In *2019 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCCom/CyberSciTech)*, pages 215–220. IEEE, 2019. [2](#)
- [29] Waqas Sultani, Chen Chen, and Mubarak Shah. Real-world anomaly detection in surveillance videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6479–6488, 2018. [3](#), [5](#), [7](#), [2](#)
- [30] Zhongyun Tang, Haiyang Hu, and Chonghuan Xu. A federated learning method for network intrusion detection. *Concurrency and Computation: Practice and Experience*, 34(10):e6812, 2022. [1](#), [2](#)
- [31] Yu Tian, Guansong Pang, Yuanhong Chen, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. *arXiv preprint arXiv:2101.10030*, 2021. [3](#), [4](#), [7](#)
- [32] Yu Tian, Guansong Pang, Yuanhong Chen, Rajvinder Singh, Johan W Verjans, and Gustavo Carneiro. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4975–4986, 2021. [3](#), [5](#)
- [33] Jhih-Ciang Wu, He-Yen Hsieh, Ding-Jie Chen, Chiou-Shann Fuh, and Tyng-Luh Liu. Self-supervised sparse representation for video anomaly detection. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, pages 729–745. Springer, 2022. [3](#), [7](#)
- [34] Peng Wu, Jing Liu, Yujia Shi, Yujia Sun, Fangtao Shao, Zhaoyang Wu, and Zhiwei Yang. Not only look, but also listen: Learning multimodal violence detection under weak supervision. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX 16*, pages 322–339. Springer, 2020. [3](#), [2](#)
- [35] Muhammad Zaigham Zaheer, Jin-ha Lee, Marcella Astrid, and Seung-Ik Lee. Old is gold: Redefining the adversarially learned one-class classifier training paradigm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14183–14193, 2020. [3](#)
- [36] Muhammad Zaigham Zaheer, Arif Mahmood, Marcella Astrid, and Seung-Ik Lee. Claws: Clustering assisted weakly supervised learning with normalcy suppression for anomalous event detection. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 358–376. Springer, 2020. [3](#), [5](#), [2](#)
- [37] Muhammad Zaigham Zaheer, Arif Mahmood, Hochul Shin, and Seung-Ik Lee. A self-reasoning framework for anomaly detection using video-level labels. *IEEE Signal Processing Letters*, 27:1705–1709, 2020. [3](#)
- [38] Muhammad Zaigham Zaheer, Arif Mahmood, Marcella Astrid, and Seung-Ik Lee. Clustering aided weakly supervised training to detect anomalous events in surveillance videos, 2022. [7](#), [2](#)
- [39] M Zaigham Zaheer, Arif Mahmood, M Haris Khan, Mattia Segu, Fisher Yu, and Seung-Ik Lee. Generative cooperative learning for unsupervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14744–14754, 2022. [2](#), [5](#), [6](#), [7](#), [8](#)
- [40] Chen Zhang, Guorong Li, Yuankai Qi, Shuhui Wang, Laiyun Qing, Qingming Huang, and Ming-Hsuan Yang. Exploiting completeness and uncertainty of pseudo labels for weakly supervised video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16271–16280, 2023. [3](#)
- [41] Hongyi Zhang, Jan Bosch, and Helena Holmström Olsson. End-to-end federated learning for autonomous driving vehicles. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021. [2](#)
- [42] Jianyi Zhang, Saeed Vahidian, Martin Kuo, Chunyuan Li, Ruiyi Zhang, Guoyin Wang, and Yiran Chen. Towards building the federated gpt: Federated instruction tuning. *arXiv preprint arXiv:2305.05644*, 2023. [1](#)